

# Spatio-Temporal Covariance Descriptors for Action and Gesture Recognition

Andres Sanin, Conrad Sanderson, Mehrtash T. Harandi, Brian C. Lovell

NICTA, PO Box 6020, St Lucia, QLD 4067, Australia \*  
University of Queensland, School of ITEE, QLD 4072, Australia

## Abstract

*We propose a new action and gesture recognition method based on spatio-temporal covariance descriptors and a weighted Riemannian locality preserving projection approach that takes into account the curved space formed by the descriptors. The weighted projection is then exploited during boosting to create a final multiclass classification algorithm that employs the most useful spatio-temporal regions. We also show how the descriptors can be computed quickly through the use of integral video representations. Experiments on the UCF sport, CK+ facial expression and Cambridge hand gesture datasets indicate superior performance of the proposed method compared to several recent state-of-the-art techniques. The proposed method is robust and does not require additional processing of the videos, such as foreground detection, interest-point detection or tracking.*

## 1. Introduction

Video-based classification plays a key role in human motion analysis fields such as action and gesture recognition. Both fields have shown promising applications in many areas, including security and surveillance, content-based video analysis, human-computer interaction and animation. According to a recent survey on recognition of human activities [28], the focus has shifted to methods that do not rely on human body models, where the information is extracted directly from the images and hence being less dependent on reliable segmentation and tracking algorithms. Such image representation methods can be categorised into global and local based approaches [22].

Methods with global image representation encode visual information as a whole. Ali and Shah [1] extract a series of kinematic features based on optical flow. A group of kinematic modes is found using principal component analysis. Guo *et al.* [6] encode the same kinematic features using sparse representation of covariance matrices. Several methods first divide the region of interest into a fixed spatial or temporal grid, extract features inside each cell and then combine them into a global representation. For example,

this can be achieved using local binary patterns (LBP) [11], or histograms of oriented gradients (HOG) [27]. Global representations are sensitive to viewpoint, noise and occlusions which may lead to unreliable classification. Furthermore, global representations depend on reliable localisation of the region of interest [22].

Local representations are designed to deal with the abovementioned issues by describing the visual information as a collection of patches, usually at the cost of increased computation. Laptev and Lindeberg [14] extract interest points using a 3D Harris corner detector and use the points for modelling the actions. One of the major drawbacks is the low number of interest points that are able to remain stable across an image sequence. A common solution is to work with windowed data, extracting salient regions which can be represented using Gabor filtering [4].

Wang *et al.* [31] showed that dense sampling approaches tend to perform better compared to interest point based approaches. Dense sampling is typically done for a set of patches inside the region of interest. Features are extracted from each patch to form a descriptor. These descriptor representations differ from grid-based global representations in that they can have an arbitrary position and size, and that the patches are not combined to form a single representation but form a set of multiple representations. Examples are HOG and HOF (histogram of oriented flow) descriptors [15], SIFT descriptors [17], and their respective spatio-temporal versions, HOG3D [31] and 3D SIFT [26]. Because of the likely large number of descriptors and/or their high dimensionality, comparing sets of descriptors is often not straightforward. This has led to compressed representations such as formulating sets of descriptors as bags-of-words [21].

In this paper we propose the use of spatio-temporal covariance descriptors for action and gesture recognition tasks. Flat region covariance descriptors were first proposed for the task of object detection and classification in images [29]. Each covariance descriptor represents the features inside an image region as a normalised covariance matrix. They have led to improved results over related descriptors such as HOG, in terms of detection performance as well as robustness to translation and scale [29]. Furthermore, covariance matrices provide a low dimensional representation which enables efficient comparison between sets of covariance descriptors.


\*Published in: IEEE Workshop on Applications of Computer Vision, pp. 103–110, 2013. <http://dx.doi.org/10.1109/WACV.2013.6475006>.

Acknowledgements: NICTA is funded by the Australian Government via the Department of Broadband, Communications and the Digital Economy, and the Australian Research Council through the ICT Centre of Excellence program.

The proposed spatio-temporal descriptors, which we name Cov3D, belong to the group of symmetric positive definite matrices which do not form a vector space. They can be formulated as a connected Riemannian manifold, and taking into account the non-linear nature of the space of the descriptors may lead to improved classification results. The most common approach for classification on manifolds is to first map the points into an appropriate Euclidean representation [16] and then use traditional machine learning methods. A recent example of mapping is the Riemannian locality preserving projection (RLPP) technique [8].

The Cov3D descriptors are extracted from spatio-temporal windows inside sample videos, with the number of possible windows being very large. As such, we use a boosting approach to search the windows to find a subset which is the most useful for classification. We propose to extend RLPP by weighting (WRLPP), in order to take into account the weights of the training samples. This weighted projection leads to a better representation of the neighbourhoods around the most critical training samples during each boosting iteration. The proposed Cov3D descriptors, in conjunction with the classification approach based on WRLPP boosting, lead to a state-of-the-art method for action and gesture recognition.

We continue the paper as follows. In Section 2 we describe the spatio-temporal covariance descriptors, and use the concept of integral video to enable fast calculation inside any spatio-temporal window. In Section 3, we first overview the concept of Riemannian manifolds formulated in the context of positive definite symmetric matrices, and then detail the proposed boosting classification approach based on weighted Riemannian locality preserving projection. In Section 4, we compare the performance of the proposed method against several recent state-of-the-art methods on three benchmark datasets. Concluding remarks and possible future directions are given in Section 5.



**Figure 1.** Conceptual demonstration for obtaining a Cov3D spatio-temporal covariance descriptor. A spatio-temporal window  $R$  is defined inside the input video. For each pixel in  $R$  a feature vector  $z_i$  is calculated. The feature vectors are then used to compute the covariance matrix  $Cov3D_R$ .

## 2. Cov3D Descriptors

In this section we first present the general form of the proposed spatio-temporal covariance descriptors (Cov3D), an algorithm for their fast calculation, and finally how they can be specialised for action and gesture recognition. For convenience, we follow the notation in [29].

Let  $V$  be the sequence of images  $\{I_t\}_{t=1}^T$  and  $F$  be the  $W \times H \times T \times d$  dimensional feature video extracted from  $V$ :

$$F(x, y, t) = \Phi(V, x, y, t) \quad (1)$$

where the function  $\Phi$  can be any mapping such as intensity, colour, gradients, or optical flow. For a given spatio-temporal window  $R \subset F$ , let  $\{z_i\}_{i=1}^S$  be the  $d$ -dimensional feature vectors inside  $R$ . The region  $R$  is represented with the  $d \times d$  covariance matrix of the feature vectors:

$$Cov3D_R = \frac{1}{S-1} \sum_{i=1}^S (z_i - \mu)(z_i - \mu)^T \quad (2)$$

where  $\mu$  is the mean of the points. Fig. 1 shows the construction of a covariance descriptor inside a spatio-temporal window. Examples of feature vectors specific for action and gesture recognition are given in Section 2.2.

Representing a spatio-temporal window with a covariance matrix has several advantages: (i) it is a low-dimensional representation which is independent on the size of the window, (ii) the impact of noisy samples is reduced through the averaging during covariance computation, (iii) it is a straightforward method of fusing correlated features.

### 2.1. Fast computation

Integral images are an intermediate image representation used for the fast calculation of region sums [30]. The concept has been extended to image sequences [10], where the integral images are stacked to form an integral video, and can be used to compute spatio-temporal region sums in constant time. For a video  $V$ , its integral video  $IV$  is defined as:

$$IV(x', y', t') = \sum_{x \leq x'} \sum_{y \leq y'} \sum_{t \leq t'} V(x, y, t) \quad (3)$$

Tuzel *et al.* [29] used the integral image representations for fast calculation of flat region covariances. Here we extend the idea for fast calculation of covariance matrices inside a spatio-temporal window using the integral video representation. The  $(i, j)$ -th element of the covariance matrix defined in (2) can be expressed as:

$$Cov3D_R(i, j) = \frac{1}{S-1} \left[ \sum_{k=1}^S z_k(i)z_k(j) - \frac{1}{S} \sum_{k=1}^S z_k(i) \sum_{k=1}^S z_k(j) \right] \quad (4)$$

where  $z_k(i)$  refers to the  $i$ -th element of the  $k$ -th vector. To find the covariance in a given spatio-temporal window  $R$ , we have to compute the sum of each feature dimension,  $z(i)_{i=1}^d$ , as well as the sum of the multiplication of any two feature dimensions,  $z(i)z(j)_{i,j=1 \dots d}$ . With  $d$  representing the number of dimensions, the covariance of any spatio-temporal window can be computed in  $O(d^2)$  time, as follows.

We need to compute a total of  $d + d^2$  integral videos. Let  $P$  be the  $W \times H \times T \times d$  tensor of the integral videos:

$$P(x', y', t', i) = \sum_{x \leq x'} \sum_{y \leq y'} \sum_{t \leq t'} F(x, y, t)(i) \quad (5)$$

where  $F(x, y, t)(i)$  is the  $i$ -th element of vector  $F(x, y, t)$ . Furthermore, let  $Q$  be the  $W \times H \times T \times d \times d$  tensor of the second-order integral videos:

$$Q(x', y', t', i, j) = \sum_{x \leq x'} \sum_{y \leq y'} \sum_{t \leq t'} F(x, y, t)(i) \cdot F(x, y, t)(j) \quad (6)$$

for  $i, j = 1, \dots, d$ . The complexity of calculating the tensors is  $O(d^2 WHT)$ . The  $d$ -dimensional feature vector  $\mathbf{p}_{x,y,t}$  and the  $d \times d$  dimensional matrix  $\mathbf{Q}_{x,y,t}$  can be obtained from the above tensors using:

$$\mathbf{p}_{x,y,t} = [P(x, y, t, 1), \dots, P(x, y, t, d)]^T \quad (7)$$

$$\mathbf{Q}_{x,y,t} = \begin{pmatrix} Q(x, y, t, 1, 1) & \dots & Q(x, y, t, 1, d) \\ \vdots & \ddots & \vdots \\ Q(x, y, t, d, 1) & \dots & Q(x, y, t, d, d) \end{pmatrix} \quad (8)$$

Let  $R(x_1, y_1, t_1; x_2, y_2, t_2)$  be the spatio-temporal window of points  $\{(x, y, t) | x_1 \leq x \leq x_2, y_1 \leq y \leq y_2, t_1 \leq t \leq t_2\}$ , as shown in Fig. 2. The covariance of the spatio-temporal window bounded by  $(0, 0, 0)$  and  $(x, y, t)$  is:

$$Cov3D_{R(0,0,0;x,y,t)} = \frac{1}{S-1} \left[ \mathbf{Q}_{x,y,t} - \frac{1}{S} \mathbf{p}_{x,y,t} \mathbf{p}_{x,y,t}^T \right] \quad (9)$$

where  $S = x \cdot y \cdot t$ . Similarly, after a few rearrangements, the covariance of the region  $R(x_1, y_1, t_1; x_2, y_2, t_2)$  can be computed as:


$$\begin{aligned} Cov3D_{R(x_1,y_1,t_1;x_2,y_2,t_2)} = & \frac{1}{S-1} \left[ \hat{\mathbf{Q}}_{x_2,y_2} + \hat{\mathbf{Q}}_{x_1-1,y_1-1} - \hat{\mathbf{Q}}_{x_2,y_1-1} - \hat{\mathbf{Q}}_{x_1-1,y_2} \right. \\ & - \frac{1}{S} (\hat{\mathbf{p}}_{x_2,y_2} + \hat{\mathbf{p}}_{x_1-1,y_1-1} - \hat{\mathbf{p}}_{x_2,y_1-1} - \hat{\mathbf{p}}_{x_1-1,y_2}) \\ & \left. (\hat{\mathbf{p}}_{x_2,y_2} + \hat{\mathbf{p}}_{x_1-1,y_1-1} - \hat{\mathbf{p}}_{x_2,y_1-1} - \hat{\mathbf{p}}_{x_1-1,y_2})^T \right] \quad (10) \end{aligned}$$

where

$$\hat{\mathbf{p}}_{x,y} = \mathbf{p}_{x,y,t_2} - \mathbf{p}_{x,y,t_1} \quad (11)$$

$$\hat{\mathbf{Q}}_{x,y} = \mathbf{Q}_{x,y,t_2} - \mathbf{Q}_{x,y,t_1} \quad (12)$$

and  $S = (x_2 - x_1 + 1) \cdot (y_2 - y_1 + 1) \cdot (t_2 - t_1 + 1)$ .



**Figure 2.** Integral feature video. The spatio-temporal window  $R$  is bounded by  $(x_1, y_1, t_1)$  and  $(x_2, y_2, t_2)$ . Each point in  $R$  is a  $d$  dimensional vector, where  $d$  is the number of features.

## 2.2. Features and regions

Commonly used features for action and gesture recognition include intensity gradients and optical flow. Previous studies have shown the benefit of combining both types of features [4, 31]. We define the feature mapping  $\Phi$ , present in (1), as the following combination of gradient and optical-flow based features, extracted from pixel location  $(x, y, t)$ :

$$\Phi(V, x, y, t) = [x \ y \ t \ g \ o]^T \quad (13)$$

where

$$g = \begin{bmatrix} |I_x| & |I_y| & |I_{xx}| & |I_{yy}| & \sqrt{I_x^2 + I_y^2} \arctan \frac{|I_y|}{|I_x|} \end{bmatrix} \quad (14)$$

$$o = \begin{bmatrix} u & v & \frac{\partial u}{\partial t} & \frac{\partial v}{\partial t} & \left( \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right) & \left( \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} \right) \end{bmatrix} \quad (15)$$


The first four gradient based features in (14) represent the first and second order intensity gradients at pixel location  $(x, y)$ . The last two gradient based features correspond to the gradient magnitude and gradient orientation. The optical-flow based features in (15) represent, in order: the horizontal and vertical components of the flow vector, the first order derivatives of the flow components with respect to  $t$ , and the spatial divergence and vorticity of the flow field as defined in [1]. Each descriptor is hence a  $15 \times 15$  matrix, as  $\Phi(V, x, y, t)$  has 15 dimensions.

For reliable recognition, several regions (and hence several descriptors) are typically used. Fig. 3 shows the spatio-temporal windows of two descriptors which can be used for recognition of face expressions. With the defined mapping, the input video  $V$  is mapped to  $F$ , a 15-dimensional feature video. Since the cardinality of the set of spatio-temporal windows  $\{R \subset F\}$  is very large, we only consider windows of a minimum size and increment their location and size by a minimum interval value. Further specifics on the windows used in the experiments are given in Section 4.

Following [29], each covariance descriptor  $Cov3D_R$ , is normalised with respect to the covariance descriptor of the region containing the full feature video,  $Cov3D_F$ , to improve the robustness against illumination variations:

$$\widehat{Cov3D}_R = \text{diag}(Cov3D_F)^{-\frac{1}{2}} Cov3D_R \text{diag}(Cov3D_F)^{-\frac{1}{2}} \quad (16)$$

where  $\text{diag}(Cov3D_F)$  is equal to  $Cov3D_F$  at the diagonal entries and the rest is set to zero.



**Figure 3.** Two examples of  $Cov3D$  windows that, together, can be useful for the recognition of face expressions.

### 3. Classification of Actions and Gestures

The Cov3D descriptors are symmetric positive definite matrices of size  $d \times d$ , which can be formulated as a connected Riemannian manifold ( $Sym_d^+$ ) [7]. In this section we first briefly overview Riemannian manifolds, followed by describing the proposed weighted Riemannian locality preserving projection (WRLPP) that allows mapping from Riemannian manifolds to Euclidean spaces. We then describe a classification algorithm that uses WRLPP.

#### 3.1. Riemannian manifolds

A manifold can be considered as a continuous surface lying in a higher dimensional Euclidean space. Formally, a manifold is a topological space which is locally similar to an Euclidean space [29]. Intuitively, the tangent space  $T_{\mathbf{X}}$  is the plane tangent to the surface of the manifold at point  $\mathbf{X}$ .

A point  $\mathbf{Y}$  on the manifold can be mapped to a vector in the tangent space  $T_{\mathbf{X}}$  using the logarithm map operator  $\log_{\mathbf{X}}$ . For  $Sym_d^+$  the logarithm map is defined as:

$$\log_{\mathbf{X}}(\mathbf{Y}) = \mathbf{X}^{\frac{1}{2}} \log\left(\mathbf{X}^{-\frac{1}{2}} \mathbf{Y} \mathbf{X}^{-\frac{1}{2}}\right) \mathbf{X}^{\frac{1}{2}} \quad (17)$$

where  $\log(\cdot)$  is the matrix logarithm operator. Given the eigenvalue decomposition of a symmetric matrix,  $\Sigma = \mathbf{U} \mathbf{D} \mathbf{U}^T$ , the matrix logarithm can be computed via:

$$\log(\Sigma) = \mathbf{U} \log(\mathbf{D}) \mathbf{U}^T \quad (18)$$

where  $\log(\mathbf{D})$  is a diagonal matrix, with each diagonal element equal to the logarithm of the corresponding element in  $\mathbf{D}$ .

The minimum length curve connecting two points on the manifold is called the geodesic, and the distance between two points is given by the length of this curve. Geodesics are related to the tangents in the tangent space. For  $Sym_d^+$ , the distance between two points on the manifold can be found via:

$$d^2(\mathbf{X}, \mathbf{Y}) = \text{trace} \left\{ \log^2 \left( \mathbf{X}^{-\frac{1}{2}} \mathbf{Y} \mathbf{X}^{-\frac{1}{2}} \right) \right\} \quad (19)$$

#### 3.2. Weighted RLPP

The usual approach for classification on manifolds is to first map the points into an appropriate Euclidean representation [29] and then use traditional machine learning methods. Points in the manifold can be mapped into a fixed tangent space (such as  $T_I$  where  $I$  is the identity matrix) [6]. Since distances in the manifold are only locally preserved in the tangent space, better results can be achieved by considering the tangent space at the Karcher mean, the point which minimises the distances among the samples, as shown in [29]. Improved results have been obtained by considering multiple tangent spaces [19, 25]. A more complex approach involves using training data to create a mapping that tries to preserve the relations between points, such as the RLPP approach [8].

RLPP is based on Laplacian eigenmaps [2]. Given  $N$  training points  $\mathbb{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N\}$  from the underlying

Riemannian manifold  $\mathcal{M}$ , the local geometrical structure of  $\mathcal{M}$  can be modelled by building an adjacency graph  $G$ . The simplest form of  $G$  is a binary graph obtained based on the nearest neighbour properties of Riemannian points: two nodes are connected by an edge if one node is among the  $k$  nearest neighbours of the other node. From the adjacency graph  $G$  we can find the degree and Laplacian matrices, respectively:

$$\mathbf{D}(i, i) = \sum_k G(i, k) \quad (20)$$

$$\mathbf{L} = \mathbf{D} - \mathbf{G} \quad (21)$$

where the degree matrix  $\mathbf{D}$  is a diagonal matrix of size  $N \times N$ , with diagonal entries indicating the the number of edges of each node in the adjacency graph.

RLPP also uses a heat pseudo-kernel matrix  $\mathbb{K}$ , with the  $(i, j)$ -th element constructed via:

$$\mathbb{K}(i, j) = k(\mathbf{X}_i, \mathbf{X}_j) = \exp \left\{ -\frac{d(\mathbf{X}_i, \mathbf{X}_j)}{\sigma} \right\} \quad (22)$$

where  $d(\cdot, \cdot)$  is the geodesic distance defined in (19).

The final mapping can be found through the following generalised eigenvalue problem [8]:

$$\mathbb{K} \mathbf{L} \mathbf{K}^T \mathbb{A} = \lambda \mathbf{K} \mathbf{D} \mathbf{K}^T \mathbb{A} \quad (23)$$

where the eigenvectors with the  $r$  smallest eigenvalues form the projection matrix  $\mathbb{A}$ .

The number of possible Cov3D descriptors inside a sample video is very large. As such, we elected to use boosting to search for a subset of the best descriptors for classification. We could use the original RLPP mapping approach to map the matrices as vectors at each boosting iteration. However, as shown in [29], the sample weights can be used to generate a mapping which is more appropriate for the critical training samples. Therefore, we propose a modified projection, specifically designed to be used during boosting, which uses sample weights to generate the final mapping. We refer to this approach as weighted Riemannian locality preserving projection (WRLPP).

In the modified projection, the adjacency graph  $G$  is replaced with a weighted adjacency graph  $\hat{G}$ , defined as:

$$\hat{G} = \mathbf{W} \mathbf{G} \mathbf{W} \quad (24)$$

where  $\mathbf{W}$  is a diagonal matrix with diagonal values that correspond to the vector of sample weights  $[w_1, w_2, \dots, w_N]$ . Using the weighted adjacency graph, edges involving critical samples (ie. samples with higher weights) become more important and their geometrical structure is better preserved. The modified projection approach is detailed in Algorithm 1.

Once the the projection matrix  $\mathbb{A}$  has been obtained, a given point  $C$  (a Cov3D matrix) on the manifold can then be mapped to Euclidean space via:

$$\text{WRLPP}(C) = \mathbf{A}^T \mathbf{K}_C \quad (25)$$

where  $\mathbf{K}_C = [k(\mathbf{X}_1, C), k(\mathbf{X}_2, C), \dots, k(\mathbf{X}_N, C)]^T$ , with  $k(\cdot, \cdot)$  defined in (22), and  $\{\mathbf{X}_i\}_{i=1}^N$  representing the training points.



---

**Algorithm 1** : obtaining weighted RLPP

---

**Input:** Training samples (covariance matrices), labels and weights  $\{(\mathbf{X}_i, y_i, w_i)\}_{i=1}^N, \mathbf{X}_i \in \mathcal{M}$

- Create Riemannian pseudo-kernel matrix:  
 $\mathbb{K}(i, j) = \exp \left\{ -\frac{d(\mathbf{X}_i, \mathbf{X}_j)}{\sigma} \right\}$  using (19) as  $d(\cdot, \cdot)$
- Construct weighted adjacency graph:  
 $\widehat{\mathbf{G}}(i, j) = \begin{cases} w_i \cdot w_j & \text{if } y_i = y_j \text{ and } \mathbf{X}_j \text{ is among the } k \text{ nearest neighbours of } \mathbf{X}_i \text{ in } \mathbb{K}. \\ 0 & \text{otherwise} \end{cases}$
- Obtain the weighted degree  $N \times N$  diagonal matrix:  
 $\widehat{\mathbf{D}}(i, i) = \sum_k \widehat{\mathbf{G}}(i, k)$
- Calculate the weighted Laplacian matrix:  
 $\widehat{\mathbf{L}} = \widehat{\mathbf{D}} - \widehat{\mathbf{G}}$
- The eigenvectors with the  $r$  smallest eigenvalues of the Rayleigh quotient  $\frac{\mathbb{K}\widehat{\mathbf{D}}\mathbb{K}^T}{\mathbb{K}\widehat{\mathbf{L}}\mathbb{K}^T}$  form the projection matrix  $\mathbb{A}$ .

**Output:** Projection model  $\lambda = \{\mathbb{A}, \{\mathbf{X}_i\}_{i=1}^N\}$

---

### 3.3. Classification

As mentioned in the preceding section, we have chosen to use boosting to find a subset of the best descriptors for classification, as the number of possible Cov3D descriptors inside a sample video is large. For simplicity, we used a combination of one-vs-one LogitBoost classifiers [5] to achieve multiclass classification.

We start with a brief description of binary LogitBoost classification, with class labels  $y_i \in \{0, 1\}$ . The probability of sample  $\mathbf{x}$  belonging to class 1 is represented by:

$$p(\mathbf{x}) = \frac{\exp\{F(\mathbf{x})\}}{\exp\{F(\mathbf{x})\} + \exp\{-F(\mathbf{x})\}} \quad (26)$$

where  $F(\mathbf{x}) = \frac{1}{2} \sum_{m=1}^M g_m(\mathbf{x})$ , with  $g(\mathbf{x})$  representing a weak learner.

The LogitBoost algorithm learns a set of  $M$  weak learners by minimising the negative binomial log likelihood of the data. A weighted least squares regression  $g_m(\mathbf{x})$  of training points  $\mathbf{x}_i \in \mathbb{R}^d$  is fitted to response values  $z_i \in \mathbb{R}$ , with weights  $w_i$ , where

$$w_i = p(\mathbf{x}_i)(1 - p(\mathbf{x}_i)) \quad (27)$$

$$z_i = \frac{y_i - p(\mathbf{x}_i)}{p(\mathbf{x}_i)(1 - p(\mathbf{x}_i))} \quad (28)$$

As we are using Cov3D descriptors (covariance matrices) as input data, we adapt the weak learners  $g_m(\cdot)$  to use the projected descriptors. In other words,  $g_m(\mathbf{x})$  is replaced with  $g_m(\text{WRLPP}(\mathbf{X}))$ , with  $\mathbf{X}$  representing a covariance matrix.

For every unique pair of classes, we train a one-vs-one LogitBoost classifier as follows. Only the samples belonging to the pair of classes are used for training the binary classifier. One class is selected to be the positive class and the other as the negative class. For each boosting iteration, we search for the region whose Cov3D descriptor best separates positive from negative samples. The descriptor is calculated for all the training samples and mapped to vector

---

**Algorithm 2** : Boosting with WRLPP

---

**Input:** Training videos with labels  $\{(V_i, y_i)\}_{i=1}^N$  belonging to  $N_c$  classes



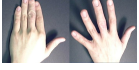
- For each unique pair of class labels  $\langle k, l \rangle$  train the one-vs-one classifier  $C_{\langle k, l \rangle}$ 
  - Let  $k$  be the positive class label and restrict the training set to  $\{(V_j, y_j)\}_{j=1 \dots N | y_j \in \{k, l\}}$
  - Let either  $k$  or  $l$  be the positive label  $y_{\langle k, l \rangle}$
  - Create binary labels  $y'_j \leftarrow (y_j = y_{\langle k, l \rangle})$
  - Start with  $w_j = 1/N, F(V) = 0, p(V_j) = \frac{1}{2}, m = 1$
  - Repeat while  $p(V_p) - p(V_n) < \text{margin}$ 
    - \* Compute the response values and weights  
 $z_j = \frac{y'_j - p(V_j)}{p(V_j)(1 - p(V_j))}, w_j = p(V_j)(1 - p(V_j))$
    - \* For each spatio-temporal window  $R_s$ 
      - Construct the descriptors  $X_{j,s} = \widehat{\text{Cov3D}}_{j,R_s}$
      - From  $\{(X_{j,s}, y'_j, w_j)\}_{j=1}^N$  obtain the projection model  $\text{WRLPP}_s$  using Algorithm 1
      - Map the data points  $\mathbf{x}_{j,s} = \text{WRLPP}_s(X_{j,s})$  using (25)
      - Fit function  $g_s(\mathbf{x})$  by weighted least-squares regression of  $z_j$  to  $\mathbf{x}_{j,s}$  using weights  $w_j$
    - \* Update  $F(V) \leftarrow F(V) + \frac{1}{2} f_m(V)$ , where  $f_m$  is the best classifier among  $\{f_s\}$  which minimises the negative binomial log-likelihood  
 $-\sum_{j=1}^N [y'_j \log(p(\mathbf{x}_j)) + (1 - y'_j) \log(1 - p(\mathbf{x}_j))]$
    - \* Update  $p(V) \leftarrow \frac{e^{F(V)}}{e^{F(V)} + e^{-F(V)}}$
    - \* Sort positive and negative samples according to descending probabilities and find samples at the decision boundaries  $V_p = (dr \cdot N_p)$ -th  $V^+$ ,  $V_n = (rr \cdot N_n)$ -th  $V^-$ , where  $dr$  and  $rr$  are the desired detection and false positive rejection rates
    - \*  $m \leftarrow m + 1$
  - Store  $C_{\langle k, l \rangle} = \{(R_m, \text{WRLPP}_m, g_m)\}_{m=1}^M$ , threshold  $\tau_{\langle k, l \rangle} = F(V_n)$  and positive label  $y_{\langle k, l \rangle}$

**Output:** A set of  $\frac{N_c(N_c-1)}{2}$  one-vs-one classifiers

---

space with WRLPP, using the sample weights calculated for the current boosting iteration. Once in vector space, we fit a linear regression and use it as the weak LogitBoost classifier.

To prevent overfitting, the number of weak classifiers on each one-vs-one classifier is controlled by a probability margin between the last accepted positive and the last rejected negative. Both margin samples are determined by the target detection rate ( $dr$ ) and the target false positive rejection rate ( $rr$ ). The final multiclass classifier is a set of one-vs-one classifiers. Each one-vs-one classifier  $C_{\langle k, l \rangle}$ , where  $k$  and  $l$  are the labels of its two classes, has a positive class  $y_{\langle k, l \rangle}$  and a threshold  $\tau_{\langle k, l \rangle}$ . The positive class is the label of the class deemed to be positive and the threshold is found via boosting. Algorithm 2 summarises the training process.

Dataset			
	UCF [24]	CK+ [18]	Cambridge [12]
Type	sports	facial expressions	hand gestures
Classes	10	7	9
Subjects	—	123	2
Scenarios	—	—	5
Video samples	150	593	900
Resolution	variable	640 × 480	320 × 240

**Table 1.** Overview of the datasets used in the experiments.

A sample video  $V$  is classified as follows. Given a one-vs-one classifier  $C_{<k,l>}$ , the probability of a sample video  $V$  belonging to the positive class  $y_{<k,l>}$  is evaluated using:

$$C_{<k,l>}(V) = \sum_{m=1}^M g_m \left( \text{WRLPP}_m(\widehat{\text{Cov3D}}_{R_m}) \right) - \tau_{<k,l>} \quad (29)$$

After evaluating  $V$  with all the one-vs-one classifiers in the set, the sample is labelled as the class  $a$  which maximises:

$$C(V) = \arg \max_a \sum_{i \neq a} C_{<a,i>}(V) \text{sign}_a(C_{<a,i>}(V)) \quad (30)$$

where  $\text{sign}_a(C_{<a,i>}(V))$  is  $\text{sign}(C_{<a,i>}(V))$  if  $a$  is the positive class  $y_{<a,i>}$ , or  $1 - \text{sign}(C_{<a,i>}(V))$  otherwise. In other words,  $V$  is labelled as the class with greater probability sum, selecting all the one-vs-one classifiers that evaluate to that class.

## 4. Experiments

We compared the performance of the proposed algorithm against baseline approaches as well as several state-of-the-art methods. We used three benchmark datasets, with an overview of the datasets shown in Table 1.


In the following subsections, we first present an evaluation of several Riemannian to Euclidean space mapping approaches, justifying the use of the weighted RLPP. We then follow with experiments showing the performance on sport actions, facial expressions and hand gestures.

Unless otherwise stated, no pre-processing was performed in the input sequences and all the recognition results were obtained using 5-fold cross validation to divide the samples into training and testing sets.

In all cases we used the following parameters: 0.95 detection rate, 0.95 false positive rejection rate, 0.5 margin. Furthermore, since the search space of spatio-temporal windows is very large, we restricted the minimum size of the windows, as well as the minimum increment on location and size of the windows, to  $\frac{1}{8}$  of the frame size.

### 4.1. Comparison of mapping approaches

In Fig. 4, we compare the following six Riemannian to Euclidean space mapping ( $\text{Sym}_d^+ \mapsto \mathbb{R}$ ) approaches which can be used during boosting: (i) no mapping (ie., using a



**Figure 4.** Performance comparison of various  $\text{Sym}_d^+ \mapsto \mathbb{R}$  mapping approaches, used within the classifier framework described in Section 3.3.

vectorised representation of the upper-triangle of the covariance matrix), (ii) projection to a fixed tangent space [6], (iii) projection to the weighted Karcher mean of the samples [29], (iv) projection using k-tangent spaces [25], (v) mapping the points with the original RLPP method [8], and (vi) mapping the points with the proposed WRLPP approach.

Since the mapping approach affects individual binary classifiers, we show results per classifier with detection error trade-off curves. We chose the one-vs-one classifiers between conflicting class pairs (where samples of one class are misclassified as the other class) on the Cambridge hand gesture recognition dataset (which is described in Section 4.4). Each point on the curve represents the average of all the chosen classifiers. The curves were obtained by varying the classification threshold  $\tau$  in Algorithm 2.

With the exception of the original RLPP method, incrementally better results are obtained by using the mapping approaches in the mentioned order, as they provide increasingly better vector representations of the manifold space. Although RLPP is designed to provide a better representation compared to tangent-based approaches, it appears not to be appropriate for boosting as it does not take into account the sample weights of critical training points. The proposed WRLPP method addresses this problem, resulting in the best overall performance.

### 4.2. UCF sport dataset

The UCF sport action dataset [24] consists of ten categories of human actions, containing videos with non-uniform backgrounds where both the camera and the subject might be moving. We use the regions of interest provided with the dataset.

We compared the Cov3D approach against the following methods: HOG3D [31], hierarchy of discriminative

Method	Performance
HOG3D [31]	85.60%
HDN [13]	87.27%
AFMKL [32]	91.30%
<b>Cov3D</b>	<b>93.91%</b>

**Table 2.** Average recognition rate on the UCF dataset [24].

space-time neighbourhood features (HDN) [13], and augmented features in conjunction with multiple kernel learning (AFMKL) [32]. HOG3D is the extension of histogram of oriented gradient descriptor [15] to the spatio-temporal case. HDN learns shapes of space-time feature neighbourhoods that are most discriminative for a given action category. The idea is to form new features composed of the neighbourhoods around the interest points in a video. AFMKL exploits appearance distribution features and spatio-temporal context features in a learning scheme for action recognition. As shown in Table 2, the proposed Cov3D-based approach achieves the highest accuracy.

### 4.3. CK+ facial expression dataset

The extended Cohn-Kanade (CK+) facial expression database [18] contains 593 sequences from 123 subjects. We used the sequences with validated emotion labels, among 7 possible emotions. The image sequences vary in duration (i.e. 10 to 60 frames) and incorporate the onset (which is also the neutral frame) to peak formation of the facial expressions.

We compared the Cov3D approach against active appearance models (AAM), constrained local models (CLM) [3], and temporal modelling of shapes (TMS) [9]. AAM is the baseline approach included with the dataset. It uses active appearance models to track the faces and extract the features, and then uses support vector machines (SVM) to classify the facial expressions. The CLM approach is an improvement on AAM, designed for better generalisation to unseen objects. The TMS approach uses latent-dynamic conditional random fields to model temporal variations within shapes.

We show the performance per emotion in Table 3, in line with existing literature. The proposed Cov3D approach achieves the highest average recognition accuracy of 92.3% (averaged over the 7 classes). The next best method (TMS) obtained an average accuracy of 87.92%.

### 4.4. Cambridge hand gesture dataset

The Cambridge hand-gesture dataset [12] consists of 900 image sequences of 9 gesture classes. Each class has 100 image sequences performed by 2 subjects, captured under 5 illuminations and 10 arbitrary motions. The 9 classes are defined by three primitive hand shapes and three primitive motions. Each sequence was recorded with a fixed cam-

Method	angry	contempt	disgust	fear	happy	sadness	surprise
AAM [18]	75.0	84.4	94.7	65.2	<b>100</b>	68.0	96.0
CLM [3]	70.1	52.4	92.5	72.1	94.2	45.9	93.6
TMS [9]	76.7	—	81.5	94.4	98.6	<b>77.2</b>	99.1
<b>Cov3D</b>	<b>94.4</b>	<b>100</b>	<b>95.5</b>	<b>90.0</b>	96.2	70.0	<b>100</b>

**Table 3.** Recognition rate (in %) on the CK+ dataset [18].

Method	Set1	Set2	Set3	Set4	Overall
TCCA [12]	81%	81%	78%	86%	82% ( $\pm 3.5$ )
PM [20]	89%	86%	89%	87%	88% ( $\pm 2.1$ )
TB [19]	<b>93%</b>	88%	90%	91%	91% ( $\pm 2.4$ )
<b>Cov3D</b>	92%	<b>94%</b>	<b>94%</b>	<b>93%</b>	<b>93%</b> ( $\pm 1.1$ )

**Table 4.** Average recognition rate on the Cambridge dataset [12].

era having roughly isolated gestures in space and time. We followed the test protocol defined in [12]. Sequences with normal illumination were considered for training while tests were performed on the remaining sequences.

The proposed method was compared against tensor canonical correlation analysis (TCCA) [12], product manifolds (PM) [20] and tangent bundles (TB) [19]. TCCA is the extension of canonical correlation analysis to multiway data arrays or tensors. Canonical correlation analysis and principal angles are standard methods for measuring the similarity between subspaces. In the PM method a tensor is characterised as a point on a product manifold and classification is performed on this space. The product manifold is created by applying a modified high order singular value decomposition on the tensors and interpreting each factorised space as a Grassmann manifold. In the TB method, video data is represented as a third order tensor and factorised using high order singular value decomposition, where each factor is projected onto a tangent space and the intrinsic distance is computed from a tangent bundle for action classification.

We report the recognition rates for the four test sets in Table 4, where the proposed Cov3D-based approach obtains the highest performance.

## 5. Conclusion

In this paper, we first extended the flat covariance descriptors proposed in [29] to spatio-temporal covariance descriptors termed Cov3D, and then showed how they can be computed quickly through the use of integral video representations.

The proposed Cov3D descriptors belong to the group of symmetric positive definite matrices, which can be formulated as a connected Riemannian manifold. Prior to classification, points on a manifold are generally mapped to an Euclidean space, through a technique such as Riemannian locality preserving projection (RLPP) [8].

The Cov3D descriptors are extracted from spatio-temporal windows inside sample videos, with the number of possible windows being very large. We used a boosting approach to find a subset which is the most useful for classification. In order to take into account the weights of the training samples, we further proposed to extend RLPP by incorporating weighting during the projection. The weighted projection (termed WRLPP) leads to a better representation of the neighbourhoods around the most critical training samples during each boosting iteration.

Combining the proposed Cov3D descriptors with the classification approach based on WRLPP boosting leads to a state-of-the-art method for action and gesture recognition. The proposed Cov3D-based method performs better than several recent approaches on three benchmark datasets for action and gesture recognition. The method is robust and does not require additional processing of the videos, such as foreground detection, interest-point detection or tracking. To our knowledge, this is the first approach proving to be equally suitable (i.e.,  $> 90\%$  recognition accuracy) for both action and gesture recognition.

Further avenues of research include adapting the method for related tasks, such as anomaly detection in surveillance videos [23], where there is often a shortage of positive examples.

## References

- [1] S. Ali and M. Shah. Human action recognition in videos using kinematic features and multiple instance learning. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 32(2):288–303, 2010.
- [2] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.
- [3] S. Chew, P. Lucey, S. Lucey, J. Saragih, J. Cohn, and S. Sridharan. Person-independent facial expression detection using constrained local models. In *IEEE International Conference on Automatic Face Gesture Recognition and Workshops*, pages 915–920, 2011.
- [4] P. Dollar, V. Rabaud, G. Cottrell, and S. Belongie. Behavior recognition via sparse spatio-temporal features. In *IEEE International Workshop on Visual Surveillance and Performance Evaluation of Tracking and Surveillance*, pages 65–72, 2005.
- [5] J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting. *Annals of Statistics*, 28(2):337–407, 2000.
- [6] K. Guo, P. Ishwar, and J. Konrad. Action recognition using sparse representation on covariance manifolds of optical flow. In *IEEE Int. Conf. Advanced Video and Signal Based Surveillance (AVSS)*, pages 188–195, 2010.
- [7] M. T. Harandi, C. Sanderson, R. Hartley, and B. C. Lovell. Sparse coding and dictionary learning for symmetric positive definite matrices: A kernel approach. In *ECCV 2012, Lecture Notes in Computer Science (LNCS), Vol. 7573*, pages 216–229, 2012.
- [8] M. T. Harandi, C. Sanderson, A. Wiliem, and B. C. Lovell. Kernel analysis over Riemannian manifolds for visual recognition of actions, pedestrians and textures. In *IEEE Workshop on the Applications of Computer Vision*, pages 433–439, 2012.
- [9] S. Jain, C. Hu, and J. Aggarwal. Facial expression recognition with temporal modeling of shapes. In *IEEE Int. Conf. Computer Vision Workshops*, pages 1642–1649, 2011.
- [10] Y. Ke, R. Sukthankar, and M. Hebert. Efficient visual event detection using volumetric features. In *IEEE Int. Conf. Computer Vision*, volume 1, pages 166–173, 2005.
- [11] V. Kellokumpu, G. Zhao, and M. Pietikainen. Human activity recognition using a dynamic texture based method. In *British Machine Vision Conference*, 2008.
- [12] T.-K. Kim and R. Cipolla. Canonical correlation analysis of video volume tensors for action categorization and detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 31(8):1415–1428, 2009.
- [13] A. Kovashka and K. Grauman. Learning a hierarchy of discriminative space-time neighborhood features for human action recognition. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 2046–2053, 2010.
- [14] I. Laptev and T. Lindeberg. Space-time interest points. In *IEEE Int. Conf. Computer Vision*, volume 1, pages 432–439, 2003.
- [15] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [16] T. Lin and H. Zha. Riemannian manifold learning. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 30(5):796–809, 2008.
- [17] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Computer Vision*, 60(2):91–110, 2004.
- [18] P. Lucey, J. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews. The extended Cohn-Kanade dataset (CK+): A complete dataset for action unit and emotion-specified expression. In *Computer Vision and Pattern Recognition Workshops*, pages 94–101, 2010.
- [19] Y. Lui. Tangent bundles on special manifolds for action recognition. *IEEE Trans. Circuits and Systems for Video Technology*, 22(6):930–942, 2012.
- [20] Y. M. Lui, J. Beveridge, and M. Kirby. Action classification on product manifolds. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 833–839, 2010.
- [21] J. Niebles, H. Wang, and L. Fei-Fei. Unsupervised learning of human action categories using spatial-temporal words. *Int. J. Computer Vision*, 79(3):299–318, 2008.
- [22] R. Poppe. A survey on vision-based human action recognition. *Image and Vision Computing*, 28(6):976–990, 2010.
- [23] V. Reddy, C. Sanderson, and B. C. Lovell. Improved anomaly detection in crowded scenes via cell-based analysis of foreground speed, size and texture. In *Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 55–61, 2011.
- [24] M. Rodriguez, J. Ahmed, and M. Shah. Action MACH a spatio-temporal maximum average correlation height filter for action recognition. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [25] A. Sanin, C. Sanderson, M. T. Harandi, and B. C. Lovell. K-tangent spaces on Riemannian manifolds for improved pedestrian detection. In *IEEE Int. Conf. Image Processing*, 2012.
- [26] P. Scovanner, S. Ali, and M. Shah. A 3-dimensional SIFT descriptor and its application to action recognition. In *International Conference on Multimedia*, pages 357–360, 2007.
- [27] C. Thurau and V. Hlavac. Pose primitive based human action recognition in videos or still images. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [28] P. Turaga, R. Chellappa, V. Subrahmanian, and O. Udrea. Machine recognition of human activities: a survey. *IEEE Trans. Circuits and Systems for Video Technology*, 18(11):1473–1488, 2008.
- [29] O. Tuzel, F. Porikli, and P. Meer. Pedestrian detection via classification on Riemannian manifolds. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 30(10):1713–1727, 2008.
- [30] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *IEEE Conf. Computer Vision and Pattern Recognition*, volume 1, pages 511–518, 2001.
- [31] H. Wang, M. M. Ullah, A. Klaser, I. Laptev, and C. Schmid. Evaluation of local spatio-temporal features for action recognition. In *British Machine Vision Conference*, 2009.
- [32] X. Wu, D. Xu, L. Duan, and J. Luo. Action recognition using context and appearance distribution features. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 489–496, 2011.